Failure-aware Lifespan Performance Analysis of Network Fabric in Modular Data Centers: Toward Deployment in Canada's North

Reza Farrahi Moghaddam,* Vahid Asghari, Fereydoun Farrahi Moghaddam and Mohamed Cheriet

Synchromedia Laboratory for Multimedia Communication in Telepresence,
École de technologie supérieure, Montreal (QC), Canada H3C 1K3
imriss@ieee.org, vahid.asghari@mail.mcgill.ca, ffarrahi@synchromedia.ca,
mohamed.cheriet@etsmtl.ca

arXiv:1401.7528v2 [cs.PF] 12 Mar 2014

Abstract

Data centers have evolved from a passive element of compute infrastructure to become an active and core part of any ICT solution. Modular data centers are a promising design approach to improve resiliency of data centers, and they can play a key role in deploying ICT infrastructure in remote and inhospitable environments with low temperatures and hydro- and wind-electric capabilities. Modular data centers can also survive even with lack of continuous physical maintenance and support. Generally, the most critical part of a data center is its network fabric that could impede the whole system even if all other components are fully functional. In this work, a complete failure analysis of modular data centers using failure models of various components including servers, switches, and links is performed using a proposed Monte-Carlo approach. This approach allows us to calculate the performance of a design along its lifespan even at the terminal stages. A class of modified Tanh-Log cumulative distribution function of failure is proposed for aforementioned components in order to achieve a better fit on the real data. In this study, the real experimental data from the *lanl05* database is used to calculate the fitting parameters of the failure cumulative distributions. For the network connectivity, various topologies, such as FatTree, BCube, MDCube, and their modified topologies are considered. The performance and also the lifespan of each topology in presence of failures in various components are studied against the topology parameters using the proposed approach. Furthermore, these topologies are compared against each other in a consistent settings in order to determine what topology could deliver a higher performance and resiliency subject to the scalability and agility requirements of a target data center design.

Keywords: Failure Analysis, Modular Data Centers, Network Topology.

1. Introduction

Traditionally, the data center (DC) term has been used as a reference to the warehouse-scale (up to tens of thousands of high-performance servers as well as the associated networking, storage and cooling equipment) facilities that centralize a company's ICT-related operations and equipment. Nowadays, DCs have become vital part of seamless operation of many daily

activities. Consequentially, the reliability and availability (uptime) of DCs and their services has been a top priority for many companies and organizations. At the same time, DCs have been gradually moving from being private properties of big players to become common commodities for all players regardless of their business scale. This move actually pushes DCs and, more generally, compute/network/content (CNC) delivery to the level of utilities.[1] The two main enablers of this move can be identified as the popular shift toward virtualization and everything as a service (XaaS)[2] and the shift toward open source paradigm at all levels of ICT operations. These enablers are getting more and more matured, and it is expected that they will become promising players in the management of the resources at large scale, such as that of DCs, in the future.

Traditional bottom-up approach in the design of DCs, which can be split into i) component (hardware and software) selection and ii) topology selection sub-designs, gives connectivity a critical role in performance and also reliability of the DC. In particular, problems such as unreachability, traffic choking at link as well as at component (switches) levels, and misconfiguration could highly affect the Return on Investment (RoI) of a design. At the same time, DC architectures and requirements can be significantly different based on the companies' and their clients' requirements. For instance, a DC designed for a cloud IaaS provider company such as Amazon requires significantly different infrastructural, reliability and security requirements compared to a DC used by a private company, such as a bank, that requires to securely maintain its client financial data and transactions.

Nonetheless, regardless of any classification for a DC, having a reliable DC network fabric with a scalable architecture design is a critical concern. By definition, in a DC network, reliability is the ability to perform and deliver the required functions and services for a specified period of time and under some conditions declared in service level agreements (SLAs). Although there are various well-known approaches to guarantee such an operation in regular DC settings, scalability in terms of expanding hardware and also installing DCs at unsecured or unmonitored places could be highly challenging. To address these challenges, DCs with a scalable architecture are designed using purpose-engineered modules that at the same time reduces costs of operation and administration in a DC. This modular data center (MDC) approach is also becoming more popular for large-scale DC operations [41]. MDCs can be imagined as set of independent, self-directed compute modules that are built in standardized sizes, such as that of overseas shipping containers, hosting a large number of servers [25].

Although software level management of DCs is a big challenge by itself, designing an agile, robust, and scalable DC hardware (topology and components) is the key element that affect

[1]The same level of electricity and water utilities.

[2]XaaS stands for IaaS, PaaS, and SaaS among others, where IaaS, for example, stands for Infrastructure-as-a-Service.

all high-level management layers. This is more critical for ship-and-deploy (S&D) MDCs that seem promising solutions for fast deployment of resilience compute resources at any place and any time. This is highly interesting, and is especially along with strategic north plans in Canada at both provincial and national levels [4, 19, 21], which target Canada's natural resources, especially untouched hydro-electric and wind-electric sites in the North [26], among other technologies. It is expected that significant infrastructure, including that of broadband telecom and ICT connectivity, will be deployed to support the development, which implicitly provide a big opportunity for ICT activities in the northern regions. Figure 1(a)-(c) shows some considerable potentials in the Canada's North in terms of natural resources and low temperature.[3] The proposed DC sites shown in Figure 1(d) reflect a potential schematic configuration in terms of hydro- or wind-electricity from the perspective of the authors. However, a detailed site analysis and feasibility study should be performed to determine the optimal number of sites and their optimal locations which is beyond the scope of this work.

Although modular designs have also become the main stream in almost all new DCs, in this study, we focus on standalone MDCs that could be possibly sent to remote areas where the Electricity, Water, and Broadband connectivity (EWBc) infrastructure is in-place, but there will be no human support or maintenance in short-term periods of time. In these scenarios and with the XaaS business perspective in mind, there would be multiple tenants who will use the low-cost, green, and ubiquitous compute resources provided by these DCs in long-term (shared ownership) or short-term (lease) service contracts [18]. The profitability of such solutions highly depends on survivability of the modules along time, and also on the quality of that survival, especially at the last mile of their terminal stage. The study of such a behavior is one of this paper's focuses.

Another aspect of the MDCs, which motivated us to consider them in this study, is their high potential in reducing energy (resource) consumption and also environmental footprint. In particular, absence of frequent human access to equipment provides an opportunity in design to effectively contain and handle the heat generated in a DC, especially using non-air (liquid) cooling system designs [5]. It has been observed that the most of electricity consumption in the traditional cooling systems is wasted to extract the heat from the hard-to-reach hot spots [3, 9, 31]. This behavior and design constraint has even resulted in some surprising conclusions, such as observing that the high inlet temperature does not a significant impact on the system-level reliability of a data center [14, 15]. However, it could be argued that a relation between the inlet temperature and reliability is not meaningful because the temperature of air at the location of chips could be in a large uncertainly interval of more than 10°C with respect to

[3]Sources:
http://atlas.nrcan.gc.ca/data/english/maps/energy/hydropower_resources_map.pdf,
http://www.windatlas.ca/en/maps.php?field=EU&height=50&season=ANU,
http://aurora.uwaterloo.ca/home/ccw/permafrost/current.

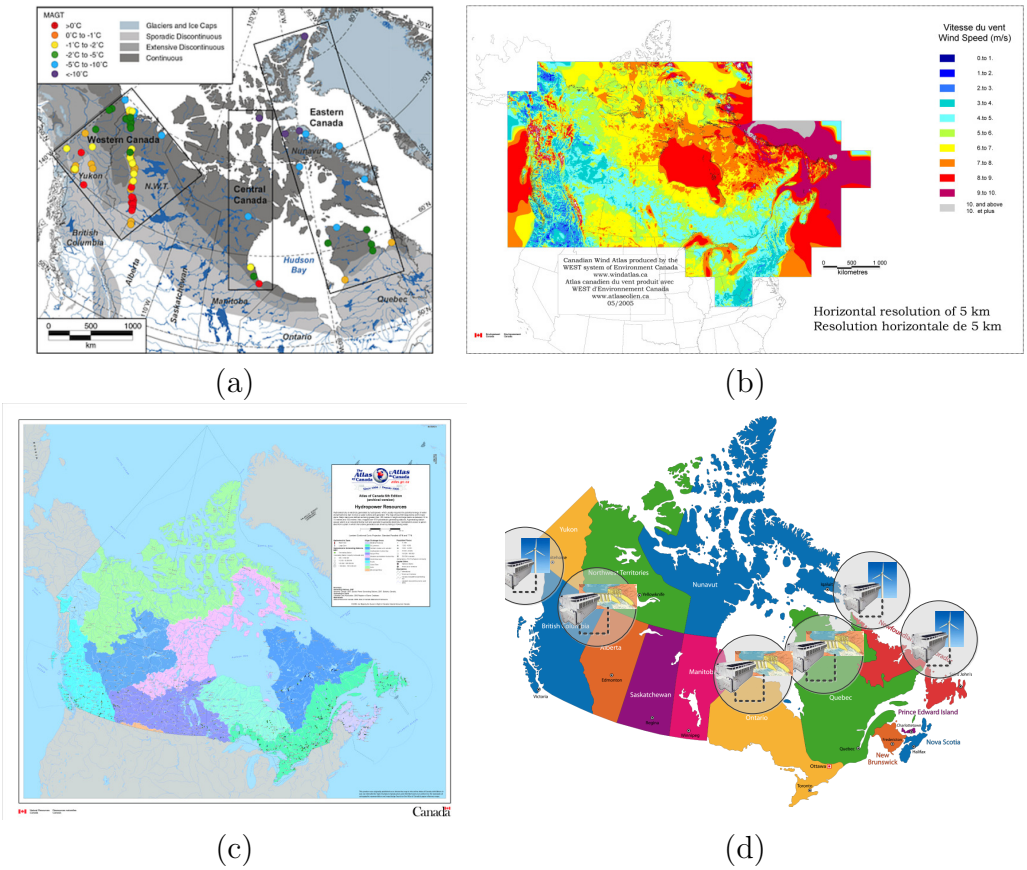

(a) (b)

(c) (d)

Figure 1: The potential of Canada's North in terms of low temperature, hydro-electricity, and wind-electricity. a) The mean annual near surface ground temperature in Permafrost of Canada's North (Source: The Canadian Cryospheric Information Network (CCIN)). b) The annually-averaged wind speed at 50m above ground of the Canadian territory (Source: Canadian Wind Energy Atlas). c) Canada's major drainage areas. There are 632 large dams and 6 major dams operational with many other in planning/deployment phases (Source: Natural Resources Canada; Earth Sciences Sector; Canada Centre for Mapping and Earth Observation). d) A proposed schematic configuration of modular, off-grid, highly efficient, and green network of DCs.

the rack inlet temperature [1, 53]. This is of great interest because the operating expenditure (OPEX) of cooling systems could highly affect the profitability of a DC. The MDCs along with distributed network of data centers [16, 17] and intelligent sensing (for example, using intelligent platform management interface (IPMI) for sensing) could drastically reduce costs, and in turn can be leveraged toward higher profit or better quality of service (QoS).

In this context, it is important to use a sophisticated methodology to compare the reliability and scalability of various DC network topologies. There are a number of work addressing the reliability concern in the DC's network, and a few DC network architectures have been investigated in order to provide a better reliability and scalability with MDCs. For example, a comparison study on the reliability of DC topologies has been performed in [23, 24] with respect to specific traffic pattern and routing algorithm, respectively. Moreover, there are a number of studies that have taken into account energy consumption and cost of various DC network architectures [40, 51]. Only recently, independent of any network applications, the

authors in [11] have investigated the reliability of DC network topologies with respect to the network components' failure ratio. It is worth noting that the reliability analysis provided in studies such as [11] were based on removing several network components according to the failure rations in order to evaluate the effect of each component type's failure on the reliability of each DC network topology. In order words, although these studies show the *relative* importance of various components, they could not provide the whole picture of degradation and decrease in reliability of the DC along time. In contrast, the operation of MDCs highly depends on such an analysis.

In this paper, we propose a new analysis approach based on individual component failure profile described in terms of cumulative distribution functions (CDFs). The proposed analysis approach, which is based on Monte-Carlo methodology, calculates both the reliability picture of the system under analysis and also several performance indicators. These indicators are proposed to provide insights in terms of profitability of a solution from various perspectives. The details are provided in the subsequent sections. In fact, up to our knowledge, there exists no study on comparing DC network topologies in terms of failures on DC components, such as servers, switches, physical links and so forth, while taking into account the operational life-time of the DC.

It is worth noting that the actual performance of a DC's network also depends on other factors. For example, the performance of network switches at their upper-bound, which is of great importance for highly-reliable designs that rely on reliability of components at their nominal ratings, may be violated if the component is designed to handle a specific performance subject to an assumed "probable" statistics. In this case, this could result in partial failure of system when the system pushes to achieve higher performance and may violate that assumed statistics. For switching components, this could be observed as choking phenomenon when all their ports are in used. Another aspect is the data flow management in the DC's network. Traditionally, switching components have been designed as Application-Specific Integrated Circuits (ASICs). This guaranteed a designated performance independent from peripheral factors, which at the same time has made them proprietary equipment. However, this vision has been gradually shifted toward more open architectures. One motivation has been the move toward Software-Defined Networks (SDNs)[4] [33] whose OpenFlow [34] is a flagship implementation for layers two and three. Other approaches to SDN implementation have been Network Functions Virtualization (NFV) [43], Programmable switches via Application Programming Interface (API), virtual network overlays and even other open virtual switch (vSwitch) approaches. In addition, routing-optimized Linux servers in the form of Open Network Installer Environment (ONIE) have been considered to completely separate hardware

[4]In the context of Software-Defined Systems (SDSs). Other examples of such a systems are Software-Defined Storage Systems (SDSSs) and Software-Defined Servers.

and software parts of switching components. This technique provides opportunities to build smart and customized switching software that enables maximum utilization of the DC's physical connectivity [35]. All these approaches are in answer to high level of under-utilization of the network connectivity in DCs, which is a problem for deadline-sensitive data transfer and flows [48, 52]. Although these flow-control mechanism are by definition topology-insensitive, the topology of the DC network could affect their performance, especially when multi-path protocols are preferred. In this work, we assume that ideal deadline-aware mechanisms are in-place at a higher level, and therefore the associated deadline failures are ignored. The effect of topology on the performance of these mechanisms will be studied in future work. In this regard, the impact of failure on the multi-path capacity of a design is studied in terms of the average max flow (Flow_{max}) metric among survived and connected nodes of every connected component (pool).[5]

In this paper, we propose a framework to study the impact of DC components' failure on the life-long operation of the whole MDC network fabric while making use of the state-of-the-art Tanh-Log characteristic function to model the failure behavior of each component type.

The paper is organized as follows. In sections 2 and 3, the three main topologies and their modifications that are considered in this study are presented and discussed. Then, in section 4, the proposed Tanh-Log cumulative distribution function is introduced and fitted using the data from a real system database. Next, the performance metrics are presented in section 5. The proposed Monte-Carlo methodology to analyze the failure performance of data centers along their lifespan is introduced in section 6. Then, the performance of the three topologies and their generalizations are studied in section 7. This is followed by the related work and some conclusions presented in sections 8 and 9, respectively.

2. Modular Data Centers (MDCs)

In this paper, we consider three typical DC network topologies found frequently in the literature, namely, FatTree [36, 37], BCube [23], and MDCube [49].[6] It is important to note that herein we do not claim that these topologies are the best, instead, we have chosen these topologies as they represent the current state-of-the-art in modular DC network topologies. This comparison study considers these topologies as a starting point for our future work on improving the understanding and analysis of failure behavior in DCs. The detailed characteristics and design of the mentioned topologies are provided below.

[5] In this work, we equivalently use the term 'pool' in place of connected component.

[6] MDCube stands for Modularized Data center Cube [49].

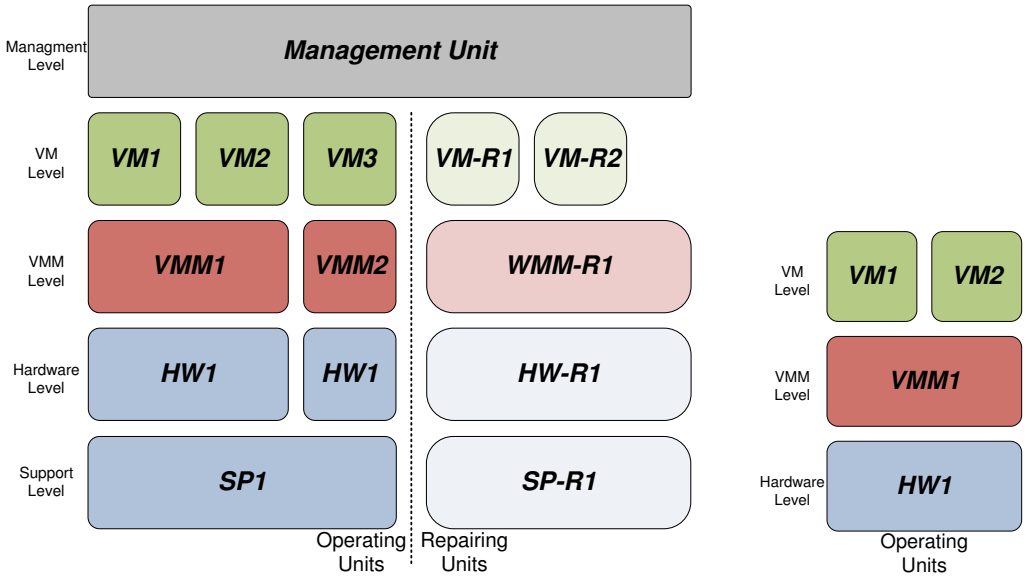

Figure 2: a) The simplified stack of components layers [18]. b) In the case of MDCs in this work that focus on the failure of the compute components, the maintenance, support and management components and layers are excluded.

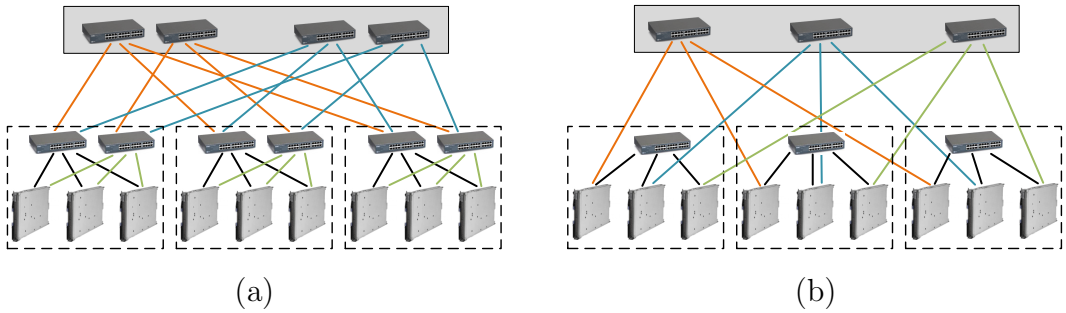

Figure 3: Two illustrative samples of a) FatTree and b) BCube topologies.

2.1. Fat-Tree Topology

The FatTree network topology is a special example of the Clos network [10], and it has been widely used as a topology for DC networks [30, 37, 39]. Herein, we refer to FatTree as the cut-down version of the classic FatTree topology proposed by Navaridas et al. in [36]. In this topology, the authors utilized the concept of FatTree, to define a DC network topology organized by reducing the number of upward ports of all switches.[7] As shown in Figure 3(a), the considered FatTree topology has two main levels, i.e., the core and the pods. The core level is composed of switches that interconnect the pods in such a way that each switch in the

[7]It is worth noting that we do not separately consider the VL2 topology proposed in [22] because of its similarity to the FatTree topology.

core is connected to at least to one switch in each pod. Each pod is composed of servers and aggregation switches. In fact, core switches are connected to the servers in each pod through the aggregation switches. Finally, each aggregation switch is connected to a different set of servers in a pod.

In this paper, we follow the K:P-ary 3-tree FatTree topology with L pods [36], which means we assume that there are L number of pods (nPod), and each pod contains K number of servers per pod (nSrvPod) with P ports on each (pSrv). Consequently, the network has $L \times K$ number of servers (nSrv), and each pod has P number of aggregation switches each with $K + P$ number of ports (pSwc). In core level, there are $2 \times P$ number of core switches each with L number of ports. Figure 3(a) shows an example of FatTree topology with $L = \text{nPod} = 3$, $K = \text{nSrvPod} = 3$ and $P = \text{pSrv} = 2$.

2.2. BCube Topology

The BCube topology was proposed as a high-performance and robust network architecture to be used to build MDCs [23]. BCube has a server-centric topology, and it provides high network capacity for various typical DC applications such as MapReduce [12]. Basically, the BCube topology is recursive and has a level-based structure. BCube topology is composed of two types of components, i.e., servers with multiple number of ports, and switches connecting a constant number of servers. In the BCube structure, multiple layers of cheap and mini Commercial-off-the-Shelf (CotS) switches could be used to connect the servers. In this topology, multiple parallel short paths are provided between any pair of servers. This not only provides high network capacity for all types of traffic, but also greatly improves redundancy, fault tolerance, and load balancing. Finally, in the BCube topology, all networking components in each level are identical, and therefore it would avoid expensive switches with high port density in higher topology levels.

A BCube topology is denoted as BCube (level-i), where i is its number of levels. A BCube (level-0) is composed of simply a single switch with K ports (pSwc) connected to K P-port servers (nSrvPod). A BCube (level-1) is constructed from N number of level-0 cubes (nPod) and N number of switches with K ports.[8] Each switch is connected to all N BCubes (level-0) networks through its connection with one server of each BCube (level-0). Similarly, the general design of a BCube (level-L) topology is composed of N number of BCubes (level-$L-1$) and N number of switches with K ports. Again, each switch of BCube (level-L) is connected with one server of each BCube (level-0). It is worth to mention that each server in a BCube (level-L) has $L + 1$ ports, which are connected to all levels from level-0 to level-L. Moreover, a BCube (level-L) has $N \times K$ total number of servers (nSrv), and $N \times (L + 1)$ number of switches (nSwc). Figure 3(b) shows the design for a BCube (level-1) with $K = \text{pSwc} = 3$,

[8]Note that K is also the number of servers in each level-0 cube (nSrvPod).

$N = \text{nPod} = 3$, $P = \text{pSrv} = 2$.

2.3. MDCube Topology

Due to the special potentials of the BCube topology, such as high network capacity and directly built from CotS switches), it is a natural step to use the BCube topology as the building block for mega-size DCs [49]. However, the BCube topology does not have enough agility, and cannot directly scaled to millions of servers just by adding more number of ports to servers and considering more number of CotS switches to the network topology. The major limitation within the design of BCubes is that the number of inter-cube cables required by BCube topology increases linearly with the increasing of the total number of servers.[9] Accordingly, the MDCube network topology has been proposed as an extension to the BCube topology by utilizing the high-speed interfaces of CotS switches to interconnect multiple containers, each container is a BCube topology by itself. From the engineering point of view, the MDCube topology design is interesting from at least two important aspects. First, it interconnects BCube containers using high-speed interfaces (10 Gbps) on the BCube's CotS switches. Note that such switches are considered commercial switches today, and having for example 48 number of 1-Gbps ports and a number of 4 10-Gbps ports is the norm. The MDCube topology uses the 10 Gbps ports for its hypercube connections. Second, the MDCube topology design can be used to create containers (clusters) of 1000-3000 servers with a high intra-connect capacity, and then interconnect them by a lower bandwidth network.[10] Note that such a designed feature is practical because there are very few application cases where there is a need for more than a few thousands servers in the same pool.

The MDCube topology can be 1-dimensional (MDCube1D) or 2-dimensional (MDCube2D). The design of MDCube topology is based on treating each BCube as a hypernode and each switch in the BCube as a hyperport. Then, a hyperport from each BCube is connected to a hyperport from another BCube in the same dimension. An example of a 2-dimensional MDCube network topology is shown in Figure 4 built from 16 BCube (level-1) containers each one with $K = \text{pSwc} = 2$ and $N = \text{nPod} = 3$.

3. The proposed modified topologies

As can be seen from the definitions of the three topologies considered in the previous section, there are some limitations in terms of the size of DCs in terms of the number of servers and also on the switches in terms of their number of ports. In particular, these topologies require various switches types with different number of ports for their different levels. At the same time, and for example for the BCube topology, the number of pods, i.e., the BCubes at level-0,

[9] It is worth to mention that we can find similar argument for the FatTree network topology.
[10] By low bandwidth we mean 10 Gbps.

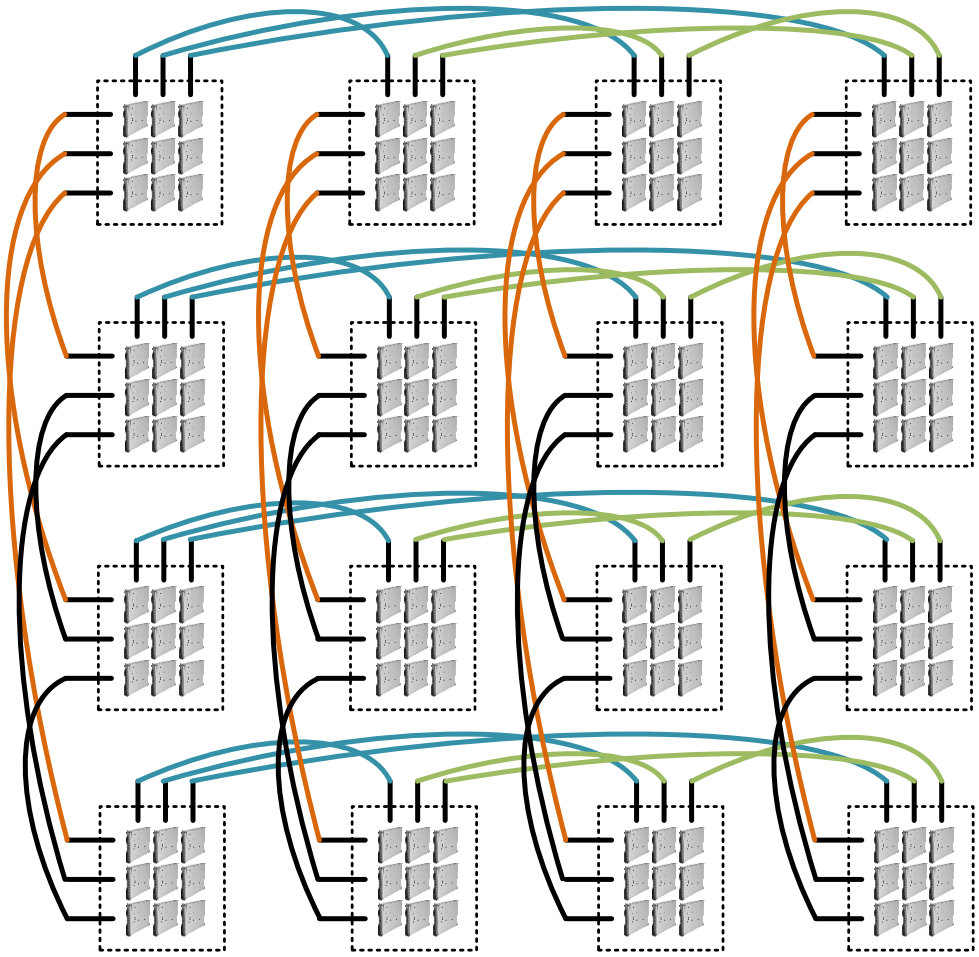

Figure 4: An illustrative sample of MDCube2D topology. There are four containers per row, each one a BCube (level-1) with 2-port servers and switches arranged in three BCube (level-0) cubes.

are restricted to the number of ports of switches. Although these constraints are bearable when the size of data center is small, they can pose considerable restrictions when several thousands of servers and switches with high number of ports are targeted. To remove these constraints, we propose extended and modified forms of these topologies in which "imaginary" switches are built using single-size actual switches. This not only reduces the complexity of the design in terms of switch port numbers, it is also helpful in making the designs more flexible in terms of the total number of servers. In the following subsections, the details of these modified topologies, which we consider then as starred version of the original topologies, are provided.

3.1. FatTree Topology*

The main restriction of the FatTree topology defined in section 2.1 is requiring high-port number switches at the level 1 and higher. This not only can impose additional requirements in the design, it also can increase the probability of large scale failure in the case of the failure of

one of main switches. In our proposed FatTree* topology, we use the same switch types across all levels. The switches' port is assigned to the underlying level in a sequential order. That means that one or more of actual switches collectively participate to create the effect of the required high port number imaginary switch compatible with the original FatTree topology. A switch may participate in one or two imaginary switches. The requirement would be to have more number of ports on the actual switches compared to the number of ports of imaginary switches. For example, consider an original FatTree topology with L pods, and K P-port servers per pod, which requires P aggregated K_P-port switches in each pod, and $2P$ core switches with L ports. The equivalent FatTree* topology would have the same configuration inside pods (K servers and 2 $K+P$-port aggregated switches). However, at the core level, we would have a $[2PL/(K+P)]$ number of $K+P$-port switches that emulate the functionality of the original high-port $2P$ switches.[11] The FatTree* topology has the advantage of using the same switching equipment in all levels. An example for $L = \text{nPod} = 3$, $K = \text{nSrvPod} = 4$, and $P = \text{pSrv} = 2$ is shown in Figure 5(a). The imaginary core switches are shown in hashed boxes. It is worth noting that the FatTree* topology is not completely equivalent to its FatTree counterpart. Especially, the fact that imaginary switches are broken on the borders of the actual switches should be considered. We will analyze and also propose modified FatTree* topologies in future to address these aspects, particularly using intra-level connectivity among switches using possibly their extra high-throughput ports.

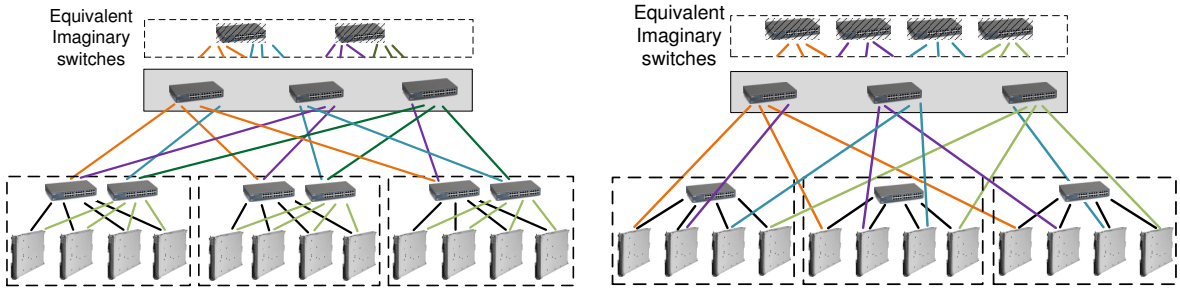

Figure 5: a) A typical FatTree* topology with $L = 3$, $K = 4$, and $P = 2$. The imaginary switches in the level-1 are shown in the dashed area. The connections shown in this area would be those if we considered the FatTree topology. b) A typical BCube* topology with $K = 4$, $N = 3$, and $P = 2$.

3.2. BCube Topology*

The same strategy, as that used to define the FatTree* topology, is used to generalize BCube to BCube*. In the proposed topology, the Level-1 switches are built using the same switch type used in the level-0 cubes. Especially, when high-port switches are used in level-0

[11] The $[\cdot]$ gives the maximum integer less than or equal to its argument.

cubes in order to scale the design, the number of these cubes is required to set to the same number of ports that is a strict restriction on the design. In contrast, we propose to choose the number of level-0 cubes independent from the number of ports of switches. Then, the level-1 switches are placed using the same switch type of the level-0 cubes with an additional condition that the sum of the number of their ports should be equal to the number of servers in all level-0 cubes. The result of this condition is that we use the same number of switches as that of the level-0 in the level-1 (and higher levels if considered). In terms of equivalency to the original BCube topology, we consider imaginary switches at the level-1 that would have the same number of ports equal to the number of cubes. Similar to the FatTree* topology, the imaginary switches of the BCube* topology are built of parts of consequential switches that are actually used in a level. An example of a BCube* topology with $K = \text{pSwc} = 4$, $N = \text{nPod} = 3$, and $P = \text{pSrv} = 2$ is shown in Figure 5(b), where K is the number of ports of switches (which is also equal to the number of servers in each level-0 cube), N is the number of level-0 cubes, and P is the number of server ports. It is worth mentioning that if original BCube topology was considered, in this example, possible configurations would be 9 (using 3-port switches) and 16 (using 4-port switches) servers. This is because of restriction to have the number of level-0 cubes equal to the port number of a level-1 switch. When, using the BCube* topology, we were able to design a network with 12 servers. This shows the flexibility of the proposed topology in terms of scale and the total number of servers. This effect is more noticeable when switches with high number of ports are considered.

3.3. MDCube2D Topology*

The MDCube2D* topology is simply defined by using the BCube* topology as the building block of the original MDCube2D topology. However, the secondary impact of this change in design would be that of high port-number switches that are easily can be used in the BCube* topology. If $K(= \text{pSwc}) < N(= \text{nPod})$ is considered for the MDCube2D*' containers (i.e., an individual BCube*), the ratio of the number of containers decreases, and therefore interconnect wiring complexity is improved. In this case, a higher number of links (wiring) among the containers would be expected. To assign the links among containers, we follow the same cyclic procedure as the original MDCube2D topology, i.e., horizontally (vertically) connecting the next available port of a container to the next to previously-connected container. Because of high number of switches per container (when $K < N$), this process should be repeated until all available ports are connected. This would lead to multiple links (multi-path possibility) between each pair of horizontal (vertical) containers. Furthermore, this redundancy in the number of links among the containers enables us to consider different degree of connectivity along horizontal or vertical directions. In other words, if we choose a less number of links per pair along the horizontal direction, for example, we would have a higher number of columns of containers along that direction. Therefore, there is more flexibility by the possibility to have

different number of columns and rows in the arrangement of the MDCube2D* topology. In contrast, we can select $K > N$ that would result in more number of containers because there will be a less number of servers per container. A single link between each pair of containers, along horizontal and vertical directions, could be expected.

4. Tanh-Log and c-Tanh-Log Failure Probability Distributions

In our approach to failure analysis, each component is assumed to follow a specific type of failure profile long time specific for its type of components. Also, we assumed that the system would be under full utilization, and therefore utilization time (age) of every surviving component is equal to the wall clock time at any moment. This implicitly imply that the system has just one cold start and there will no (hardware) repair in the life-time of the system before major maintenance cycles. This is compatible with our MDC use cases in which the systems are shipped, deployed, and remotely controlled until their RoI goes beyond a specific value that triggers a major maintenance cycle.

In Figure 6, a typical failure profile, which is actually the associated failure CDF of a typical component, is shown. The key behaviors to highlight are burn-in period at the beginning of utilization followed by a *sigmoid*-like convergence to most-likely failed state. The burn-in period is illustrated as the time interval between the blue and green vertical lines. Although detailed analysis of such behaviors is not the focus of this paper, there are some interesting points to mention.

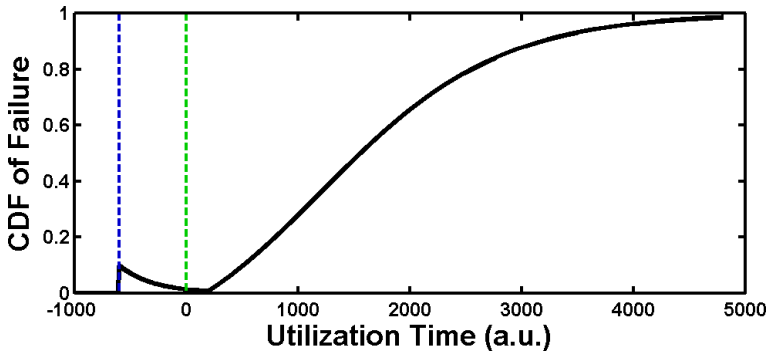

Figure 6: A typical failure profile of a component or system considering the burn-in period. The blue vertical line marks the manufacturing time of the component, while the green vertical line stands for the beginning of its utilization by the user.

First, as can be seen from the figure, by adjusting the end of the burn-in period, it is possible to deliver a component at its lowest likely-to-fail state. Although this is of great interest to high performance computing (HPC) applications and is actually practiced, it implicitly implies that a large number of manufactured, working components would not be delivered and are aborted within the burn-in phase [7, 46, 47, 50]. Despite the fact that this practice increases the reliability of the target HPC applications, it avoids the aborted components

from entering into the market while otherwise they could well serve in other applications with low-sensitivity to reliability. This aspect will get more attention in future because it is being argued that the sustainability audit of a system should not only include its immediate action but also the consequences of its indirect actions, such as its influence in the manufacturing and recycling (end of the life) of even those components that have not been used within that system boundary. This approach to the life cycle assessment (LCA) is getting more and more attention in recent years [13, 38], and it requires proper consideration in system design and operation [42]. In particular, non-invasive approaches to separate high-performance and low-performance manufactured components could be an essential replacement of the brute-force burn-in tests, which also would being a new stream of revenue and creditably for the 'responsible' manufacturers.

Second, although we can conclude a sigmoid-like reliability behavior for a system similar to those of its components, the system availability cannot be bounded with only this information. In practice, the system will go through several unavailability episodes during its lifetime because of repairable failures at higher levels of the system. A simplified schematic of the system is shown in Figure 2(a) [18]. Aging phenomena in the software levels, such as Virtual Machine Managers (VMMs) level, would eventually lead the system, and also the individual virtual infrastructures hosted on it, to fail or degrade to an unacceptable level of Service Level Objectives (SLOs). These episodes could be treated with manual or automatic interventions that can be easily performed remotely. Analysis of the (sub)-system availability and the impact of individual components, such as VMMs, is beyond the scope of this work. However, the lifetime analysis provided herein could be of great interest in planning high availability designs, since the availability of a system cannot go beyond its reliability lifetime.

In failure analysis of systems with aging and history phenomena, proper selection of failure probability distribution is a critical step. A new Tanh-Log cumulative distribution function was introduced in [18] with better fitting features. It provides flexibility to absorb various behaviors thank to its Tanh shape and also nonlinear transformation of the time. The original Tanh-Log CDF, provide in equation (1), assumes a zero failure rate at utilization time equal to zero, which is not completely consistent with the fact that the failure rate at the end of the burn-in tests could be non-zero.

$$\mathrm{CDF}^{x_c,x_s}_{\text{Tanh-Log}}(x) = \frac{1}{\tanh\left(\frac{x_c}{2\,x_s}\right)+1}\left\{\tanh\left(\frac{x_c}{2\,x_s}\right)+\tanh\left(\frac{\log(x)-x_c}{2\,x_s}\right)\right\}, \tag{1}$$

$$DDF^{x_c,x_s,s}_{\text{Tanh-Log}}(C^*) = \frac{s - s\tanh\left(\log\left(\mathrm{e}^{\frac{x_c}{2\,x_s}}\sqrt{-\frac{C^*-1}{C^*\,\mathrm{e}^{\frac{x_c}{x_s}}+1}}\right)\right)^2}{2\,x_s + 2\,x_s\,\tanh\left(\frac{x_c}{2\,x_s}\right)}, \tag{2}$$

where $C^* = \text{CDF}^{x_c,x_s}(x^*)$ at a typical time interval x^*, and DDF stands for the Differential Distribution Function, defined as follows, and provides a mean to compare even those CDFs defined on different time intervals [18]:

$$\text{DDF}^s\Big(C^* = \text{CDF}(x^*)\Big) := \frac{1}{s}\frac{\partial \text{CDF}(x)}{\partial \log(x)}\bigg|_{x^*}$$

where s is a scaling factor.

In order to remove the limitation of having a zero CDF value at $t = 0$, we use a modified version of this distribution, called the c-Tanh-Log distribution, and provided below:

$$\text{CDF}^{x_c,x_s}_{\text{c-Tanh-Log}}(x) = c_0 + \frac{1-c_0}{\tanh\left(\frac{x_c}{2\,x_s}\right)+1}\left\{\tanh\left(\frac{x_c}{2\,x_s}\right) + \tanh\left(\frac{\log(x)-x_c}{2\,x_s}\right)\right\}, \tag{3}$$

$$DDF^{x_c,x_s,s}_{\text{c-Tanh-Log}}(C^*) = \frac{s\left(\tanh\left(\log\left(\mathrm{e}^{\frac{x_c}{2\,x_s}}\sqrt{\frac{C^*-1}{c_0+c_0\,\mathrm{e}^{\frac{x_c}{x_s}}-C^*\,\mathrm{e}^{\frac{x_c}{x_s}}-1}}\right)\right)^2-1\right)(c_0-1)}{2\,x_s\left(\tanh\left(\frac{x_c}{2\,x_s}\right)+1\right)}, \tag{4}$$

where c_0 is a constant to absorb the non-zero behavior at the beginning of the utilization phase. To show the performance of the new distribution, the fit of the two distributions on the empirical *lanl05* database [44] is shown in Figure 8. The time interval of the *lanl05* database is considered to be generic, and therefore it is considered with an arbitrary unit (a.u.) of time. The biggest interval in this database is 34, 480 a.u. of time, which determines the interval on all the best fit figures. The empirical CDF of the (*union*-interpreted) *lanl05* database was retrieved from the Failure Trace Archive (FTA) [28]. As can be seen from the figure, the c-Tanh-Log distribution provides a better fit at the beginning. The fitted parameters of the two distributions are $x_c = 5.564(\pm 0.0035)$ and $x_s = 1.577(\pm 0.0030)$, and $x_c = 5.643(\pm 0.0019)$, $x_s = 1.536(\pm 0.0012)$, and $c_0 = 0.01755(\pm 0.0004)$ respectively. For each parameter value, the %95 interval is shown using $(\pm\cdot)$ notation. The best fits of both distributions and also that of the well-known Weibull distribution [28][12] are shown in Figures 8(a) and 8(b). The empirical CDF is shown in black, while the Tanh-Log and the Weibull distributions are shown in solid blue and dashed red lines respectively. The p-values of fit (compared to the traditional significance level of 0.05) with respect to the Kolmogorov-Smirnov [27, 32] and also the Anderson-Darling [45] goodness of fit (GOF) tests are 0.4765 and 0.5659,

[12]Among various CDF functions that have been used in the literature, the Weibull distribution is considered in this study as the baseline because of its effectiveness for large-scale systems [29]. This distribution has two parameters: the shape parameter β and the scale parameter δ: $\text{CDF}^{\beta,\delta}_{\text{Weibull}}(x) = 1 - \exp^{-(x/\delta)^\beta}$.

and 0.4969 and 0.5898 for the two distributions respectively.[13] In this work, we use this empirical distribution as our model for the servers. For the switches and network links, we assume the performance is higher, and therefore availability intervals of the same distribution multiplied by a factor of 2 and 5 are considered respectively. The same process of fitting are performed on these intervals in order to calculate the c-Tanh-Log distributions of these components. The best parameters of all three types of components are listed below:

$$\begin{array}{rcl} x_c^{srv} & = & 5.643\ (\pm 0.0019),\ x_s^{srv} = 1.536\ (\pm 0.0012),\ \text{and}\ c_0^{srv} = 0.01755\ (\pm 0.0004) \\ x_c^{swc} & = & 6.337\ (\pm 0.0023),\ x_s^{swc} = 1.536\ (\pm 0.0014),\ \text{and}\ c_0^{swc} = 0.00864\ (\pm 0.0005) \\ x_c^{lnk} & = & 7.253\ (\pm 0.0029),\ x_s^{lnk} = 1.536\ (\pm 0.0017),\ \text{and}\ c_0^{lnk} = 0.00152\ (\pm 0.0007) \end{array} \quad (5)$$

It is worth noting that the proposed failure analysis, described in section 6, is completely independent from the distributions that are considered for the components, and the results can be recalculated when empirical distribution for any specific component become available. The best fits of the Weibull and c-Tanh-Log distributions for the switch and link components are shown in Figure 9. The time interval is limited to that of servers because after that point there will be no connected component (pool) of servers available. However, the distributions associated to switches and links theoretically go beyond that point till they converge to a value of one in terms of the CDF values.

Because of the MDCs nature, it is assumed that all components are fully maintained/repaired to their best status at $t = 0$. To analyze the performance of an MDC, we calculate various performance metrics on its *snapshots*. A snapshot of a system is defined as a probable instance of that system with some of its components failed according to their CDFs. Every CDF is a function of time, therefore the snapshots taken at early stages of utilization are highly different from those taken in final stages. In our study, the behavior of the system's path (trajectory) to failure is not considered mostly because of sealed-box nature of the MDCs. The sealed-box assumption makes the changes in the system's state mostly directed (and toward the total failure state) without observing cycles or other complex sub-trajectories, which are more common in the open designs. The performance metrics are first introduced and reviewed in the next section, and then a methodology is introduced in order to generate and analysis the snapshots of an MDC system, and calculate its performance against the proposed metrics.

[13] To calculate the p-values, averaging over 1000 p-value estimations, each of which was calculated on a randomly-selected set of 30 samples from the real data set, was used. To perform this process, the FTA toolbox [28] was used.

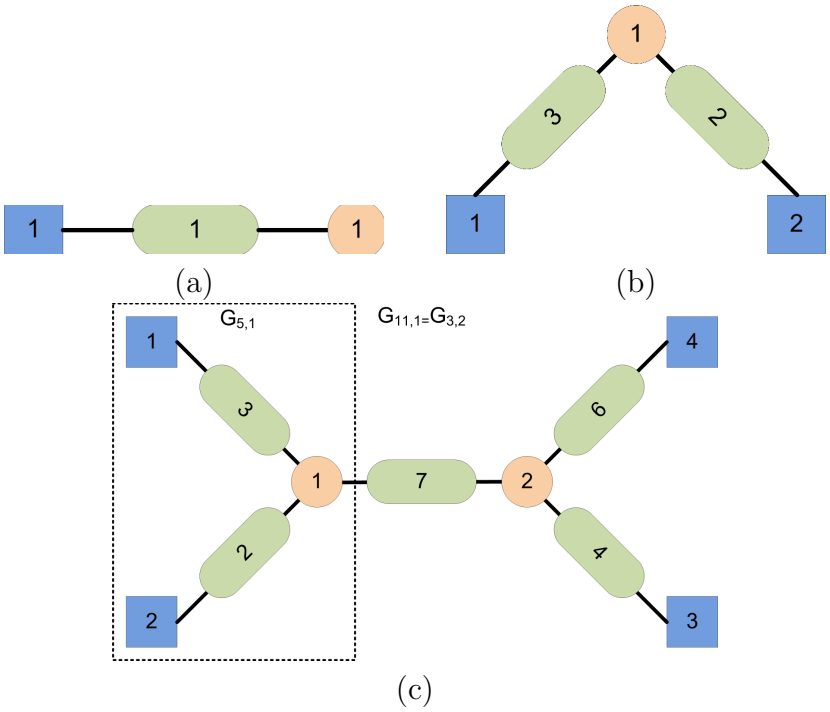

Figure 7: Various examples of sub-graphs.

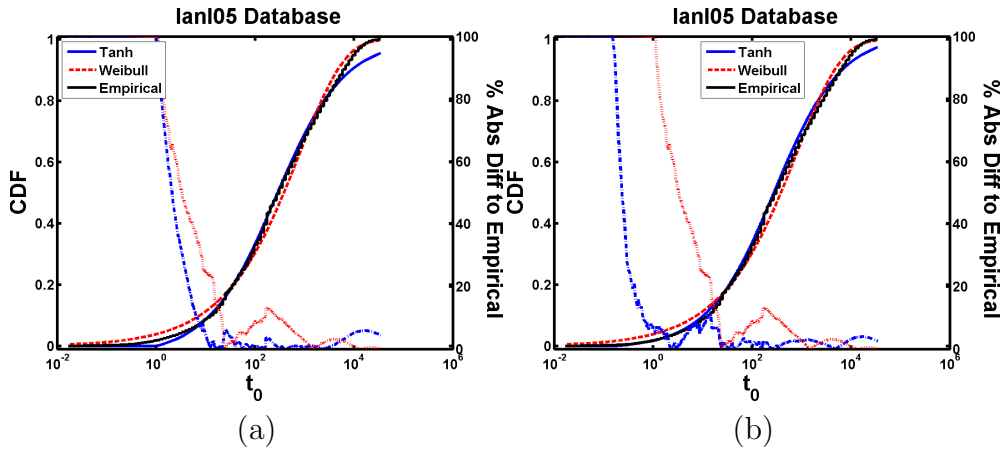

Figure 8: a) The empirical CDF of the *lanl05* database compared with its best Weibull and Tanh-Log fits. b) The same as (a) but with the c-Tanh-Log distribution. The relative difference percentage between the empirical distribution and the best fits are also shown in each figure.

5. Performance Metrics

As mentioned before, in our analysis, we only consider the failure behavior of the main components in a DC network topology, i.e., servers (blades), switches and physical links. However, considering data-centric applications, the system performance is mainly dependent on the server components. Therefore, in definition of various size-related metrics below, only the number of servers in each pool is considered. The switches and links play role in sustainability of the pools and also the maximum flow achievable each two servers. All metrics in this section

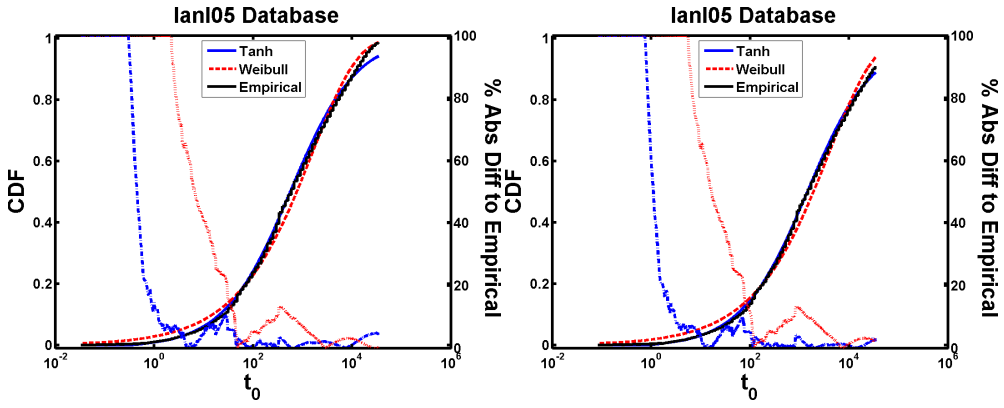

Figure 9: The tanh fits for a) switches and b) network connections.

are defined on a system snapshot, i.e., an instance of the system with some of components failed. The process to generate the snapshots and also the procedure to aggregate the metrics calculated on the set of snapshots are discussed in the next section.

The metrics related to the pool size in each snapshot are the maximum and average relative sizes of the connected components after failures in DC devices, which are defined in analogy to the metrics used in a study on reliability in complex networks [2, 20]. The relative size of the largest connected component to the total number of servers in all pools, called Maximum relative size and denoted $\mathrm{RS}_{\max}$, is given by Equation (6):

$$\begin{aligned} RS_{\max}(t_1) &= \max_{i \in I} \{RS(t_1, i)\} \\ &= \frac{\max_{i \in I} \{n_{\mathrm{srv},i}(t_1)\}}{\sum_{i \in I} n_{\mathrm{srv},i}(t_1)} \end{aligned} \tag{6}$$

where $n_{\mathrm{srv},i}(t_1)$ denotes the number of servers available in the i^{th} connected components at time t_1. It is worth mentioning that although various snapshots are possible for any time interval t_1, we use only the time value to index the snapshots in the metrics for the purpose of simplicity of notation.

5.1. Maximum absolute relative size ($ARS_{\max}$)

Absolute relative size (ARS) of a connected component is defined as the ratio of the number of servers in that connected component to the total number of servers of initial state. Then, the Maximum absolute relative size, $\mathrm{ARS}_{\max}$, is defined as the maximum of all individual ARS values of all system's connected components at time t_1:

$$\begin{aligned} \mathrm{ARS}_{\max}(t_1) &= \max_{i \in I} \{ARS(t_1, i)\} \\ &= \frac{\max_{i \in I} \{n_{\mathrm{srv},i}(t_1)\}}{N_{\mathrm{srv}}} \end{aligned} \tag{7}$$

where N_{srv} denotes the total number of servers in all the connected components at time t_0, i.e., $N_{\mathrm{srv}} = \sum_{i \in I} n_{\mathrm{srv},i}(t_0)$, where t_0 denotes the starting time of the MDC utilization. In this work, $t_0 = 0$.

5.2. Total absolute relative size (ARS_{total})

For a single leaser (tenant) of the MDC, the maximum absolute relative size ($\mathrm{ARS}_{\max}$) provides the maximum service level that can be provided to that leaser (for example, an enterprise). However, DCs are most likely operate in *multi-enterprise*[14] configurations, and therefore, the total profitability and revenue of the DC cannot be estimated from $\mathrm{ARS}_{\max}$. To address this need, we introduce the Total absolute relative size metric, $\mathrm{ARS}_{\mathrm{total}}$, that counts the total number of surviving servers in the *big-enough* pools to the initial number of servers. To define the profitable pools, a minimum threshold for the the number of servers of a pool is considered, and all pools with less that this threshold size are ignored in the calculations of $\mathrm{ARS}_{\mathrm{total}}$. The minimum threshold is denoted $n_{\mathrm{srv,req}}$:

$$\mathrm{ARS}_{\mathrm{total}}(t_1, n_{\mathrm{srv,req}}) = \frac{\sum_{i \in I} \left[\frac{n_{\mathrm{srv},i}(t_1)}{n_{\mathrm{srv,req}}} \right] \times n_{\mathrm{srv,req}}}{N_{\mathrm{srv}}} \tag{8}$$

where $[\cdot]$ gives the maximum integer less than or equal to its argument.

5.3. Maximum network flow ($Flow_{\max}$)

Considering increasing demand to process bigger volumes of data and also physical limits of maximum data transfer rate on copper or optical links, multi-path data transfer has attracted a considerable attention, and with ongoing efforts on SDNs and other initiatives to ease the restrictions of software components, the multi-path data transfer in DCs would reach the hardware upper limit of multi-path transferring in near future. Here, we introduce a metric to evaluate the quality of a topology in terms of multi-path data transfer considering failure in its components. The maximum network flow, $\mathrm{Flow}_{\max}$, of the i^{th} connected components at the time t_1 is defined as follows:

[14] Also, it is informally spelled as multi-tenant.

$$\text{Flow}_{\max}(t_1, i) = \text{mean}\left(\left\{\text{Flow}(\text{srv}_l, \text{srv}_m) \;\middle|\; \forall\, l, m \in i\right\}\right) \tag{9}$$

where $\text{Flow}(\text{srv}_l, \text{srv}_m)$ denotes the flow capacity between two of the available servers in the i^{th} connected components. In the experimental results, we will study the $\text{Flow}_{\max}$ of the largest and second-largest pools. In the next section, the proposed analysis method that generates, evaluates, and aggregates snapshots to create the performance picture of a DC and its associated topology is discussed.

6. Monte-Carlo (MC) Analysis of Topologies

Let's consider an MDC with specified topology and components. It is assumed that the system is fully operational at $t = 0$. Therefore, there is only one state of the system is possible at time $t = 0$. For any $t > 0$, the set of reachable system's states is equal to the set of all possible system's states. However, the probability of having any of those states could be different depending on t and the degree of failure in the system. In the proposed MC approach, a set of snapshots for each t is taken from its set of possible system's states according to their probabilities. These sets of snapshots along the time provide a statistical picture of the performance of the specific topology of the MDC measured using the performance metrics aggregated on the snapshots.

In order to make the MC process more direct, the snapshots are directly created from the states of individual components (servers, switches, and network links) at a specific t instead of choosing a snapshot configuration and then calculating its probability. In the direct approach, for each component, the CDF of failure is used along our in-house random number generator in order to determine its state in terms of being operational or failed. This process is repeated for all the components, and at the end a snapshot of the system at that specified instance of time is created. The whole process of creating a snapshot is repeated to have a statistically-meaningful set of snapshots for each time instance t. We have used a size of 100 snapshots per time instance in this study. A schematic diagram illustrating the concepts of paths and snapshots is provide in Figure 10. Each graph represents a snapshot of a MDC at its associated time. Only two of many possible trajectories are drawn. At $t = 0$, all trajectories have the same snapshot, from which they diverge to their own path. The MDC used in the figure is a MDCube2D* topology with 256 servers and 128 switches. The servers are represented by blue nodes, while red nodes correspond to the switches. There is a convergence at the end of life along all trajectories and paths. However, it is not in our interest because the performance of the system is negligible at that period, and the maintenance event should be schedule before. It is worth noting that in the proposed analysis approach, as mentioned before, the snapshots of a time instance t are not built based on the snapshots taken from its preceding instance

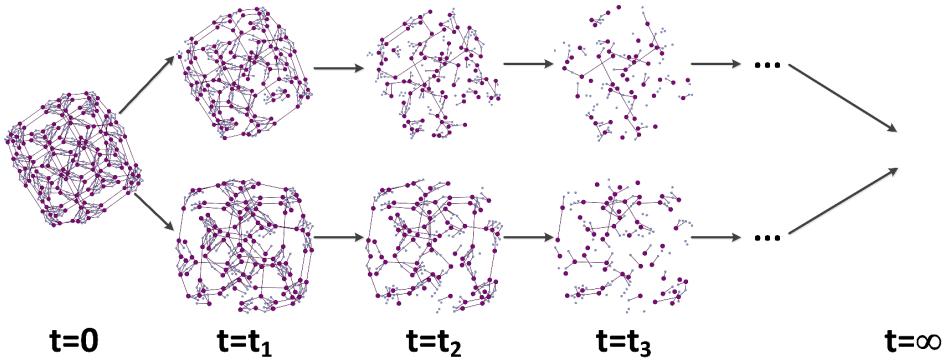

Figure 10: A typical and illustrative example of evolution of an MDC topology along the time in presence of failure. The MDC topology is an MDCube2D* with 256 servers and 128 switches, which are shown in blue and red nodes, respectively. The system can fall in one of many possible trajectories two of them are shown in the figure.

t^-.[15] Although it seems that following the states of time instances could give a more coherent picture of the system behavior, it will introduce bias in the calculations by skewing them toward those trajectories that have been selected by the random numbers used to generate the snapshots at the beginning of the utilization. The size of MC sampling set is limited (100 in our case), and therefore the aforementioned bias is unavoidable if trajectories are followed.[16] To mitigate this phenomenon, the snapshots of each time instance are independently generated from those snapshots have been selected for the receding time instances. Also, there is no special case considered for various types of components because in reality it is impossible to operate a system while forcing no failure for a specified type of components. Although such an analysis would give some insights on what component type could be improved in terms of failure resiliency in order to increase the performance, the nature of MDCs, and especially being network-centric, puts switch and link components on the higher required levels of resiliency. This fact has been considered in the CDFs discussed in section 4.

To analyze a specific topology with specified parameters (that determine its size, level of redundancy, and connectivity), the snapshots of each time instance are aggregated to calculate all performance metrics and their confidence intervals. Then, the picture of the system behavior along the time is built by concatenating these metrics. In the next section, the performance of the three topologies considered in this study is compared and evaluated against various parameters.

[15] t^- is an indistinguishable time instance *before* t.

[16] A sample video of the evolution of a MDCube2D* topology along one of its trajectories can be viewed here: http://youtu.be/Y0G7xEiY4fk.

7. Results and Discussions

In this section, the performance of each topology against its parameters is evaluated in the following three subsections.

7.1. BCube topology

As mentioned in section 2.2, the BCube topology has two parameters, N which is the number of ports of each server (pSrv), and K which is the number of ports of each switch (pSwc). By definition, N is also the number of switches in each level of the BCube, and the number of levels, and K is the number servers in each BCube (level-0), and the number of the BCube (level-0)s in the zeroth level. The performance of BCubes against these two parameters is provided in Figures 11. In Figure 11(a), the performance of the ARS_{max} against the number of switches' port is illustrated. The number of servers' port is set fixed to 5. As can be seen with the increase in the number of ports, the behavior converge to a smooth and saturated form. The gain achieved at the inflection time, i.e., the time the performance curve changes its curvature behavior,[17] is almost %100 in this figure. In the rest of plots in Figure 11, the inflection time is only provided in order to avoid plotting many curves. As can be seen, most of the performance parameters have improved behavior when the number of ports is increased. A comparison with other topologies and with higher and more practical number of servers and switches is provided in subsection 7.4.

7.2. FatTree topology

The FatTree topology is considered as the baseline in this work. The performance metrics behavior against parameters of this topology is provided in Figure 12. The main parameters of FatTree topology, defined in section 2.1, are the number of servers' port (pSrv), the number of pods (nPod), and the number of servers per pod (nSerPod). When comparing Figure 12(a) and 11(a), it can be seen that the saturated performance of the FatTree topology is achieved by the BCube topology. The main cost associated with the BCube topology is bifunctional role of servers in this topology where they participate in data transmission among other nodes. A possible improvement could be designing isolated micro switches that reside within the physical box of a server and could possibly share some of non-compute resources, such as cooling, electricity power, and even board with the host server. This not only reduces the chance of performance degradation or resource-based attacks on the server [6], it also improves the overall efficiency in terms of spatial and other resources-usage perspectives. We will study this concept in the future.

[17] The inflection time is calculated using the Extremum Distance Estimator (EDE) method [8].

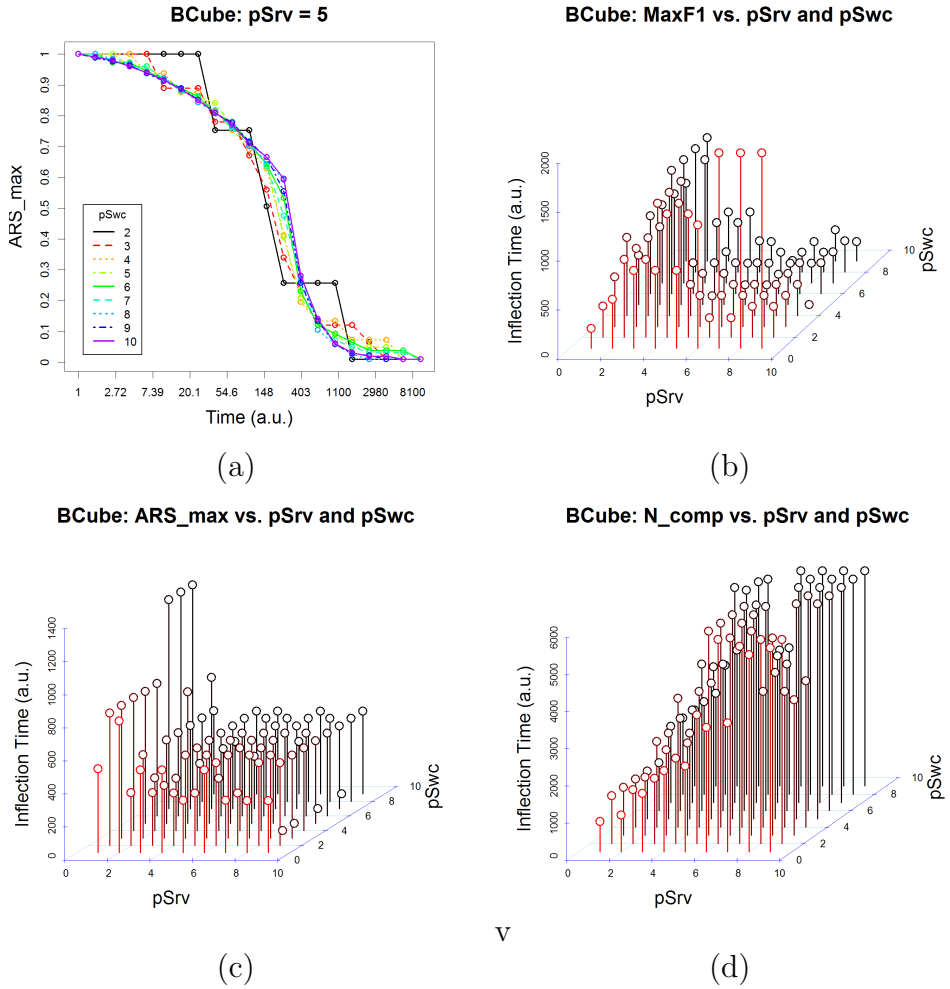

Figure 11: a) The profile of the ARS_{max} metric for the BCube topology. b)-d) The variation of $Flow_{max}$, ARS_{max}, and N_{comp} metrics versus the BCube topology's parameters.

7.3. MDCube2D topology

The same process of evaluating performance metrics against the parameters of the MDCube2D topology were carried out and is provided in Figure 13. Obviously, the MDCube2D topology has less performance in terms of ARS_{max} as can be seen from Figure 13(a). This can be associated with the bottlenecks created by the interconnecting links among the containers. However, the behavior of this topology is more stable across the parameter values, which can be seen from Figure 13(c). The other advantage of the MDCube2D topology is its by-definition scalability. When the BCube topology is used for a high number of servers, the number of level-0 cubes and also the complexity of the connections rapidly increase. In contrast, MDCube2D topology is highly scalable in a 2D manner, which fits well with the spatial organization of the traditional DC rooms.

The compromise between scalability and performance can be invested on depending on the application context. For example, for a central data center, which requires adding new modules

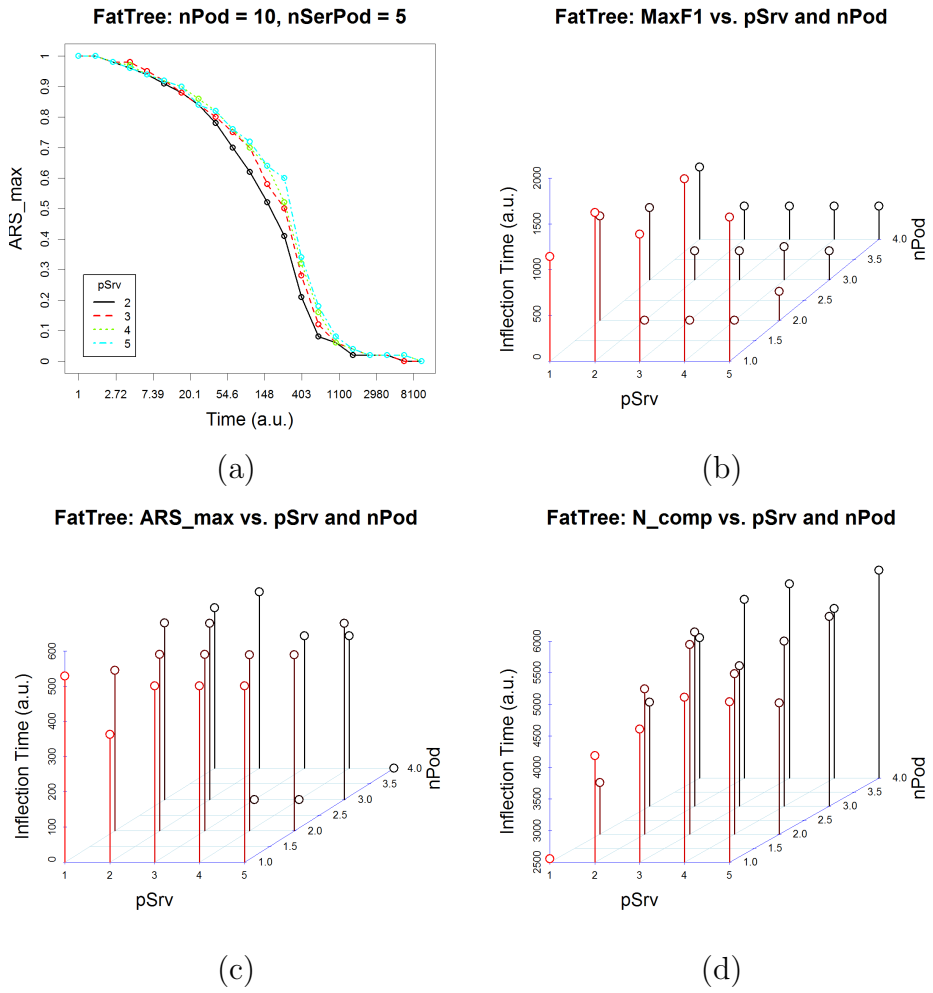

Figure 12: a) The profile of the ARS_{max} metric for the FatTree topology. b)-d) The variation of $Flow_{max}$, ARS_{max}, and N_{comp} metrics versus the FatTree topology's parameters.

and hardware along the operating in order to address the needs of new customers, it can be suggested that the MDCube2D topology is a better choice in terms of scalability. In contrast, for example, in a small edge data center used by a Telco operator to service its customers in a small area, the BCube topology seems to be the better choice in order to achieve longer maintenance-free operating intervals and also to reduce the costs associated with downtimes and maintenance. In the next section, section 7.4, we compare the performance of the extended topologies proposed in this paper are compared against each other in the practical scales of a few thousands of servers.

7.4. A use case of 3072-server MDC

All three considered topologies show a considerable resilience to failure of their components along time. However, as observed in the previous sections, the degree of these residencies varies among the topologies. Also, the comparison at the scales similar to real word DCs should be

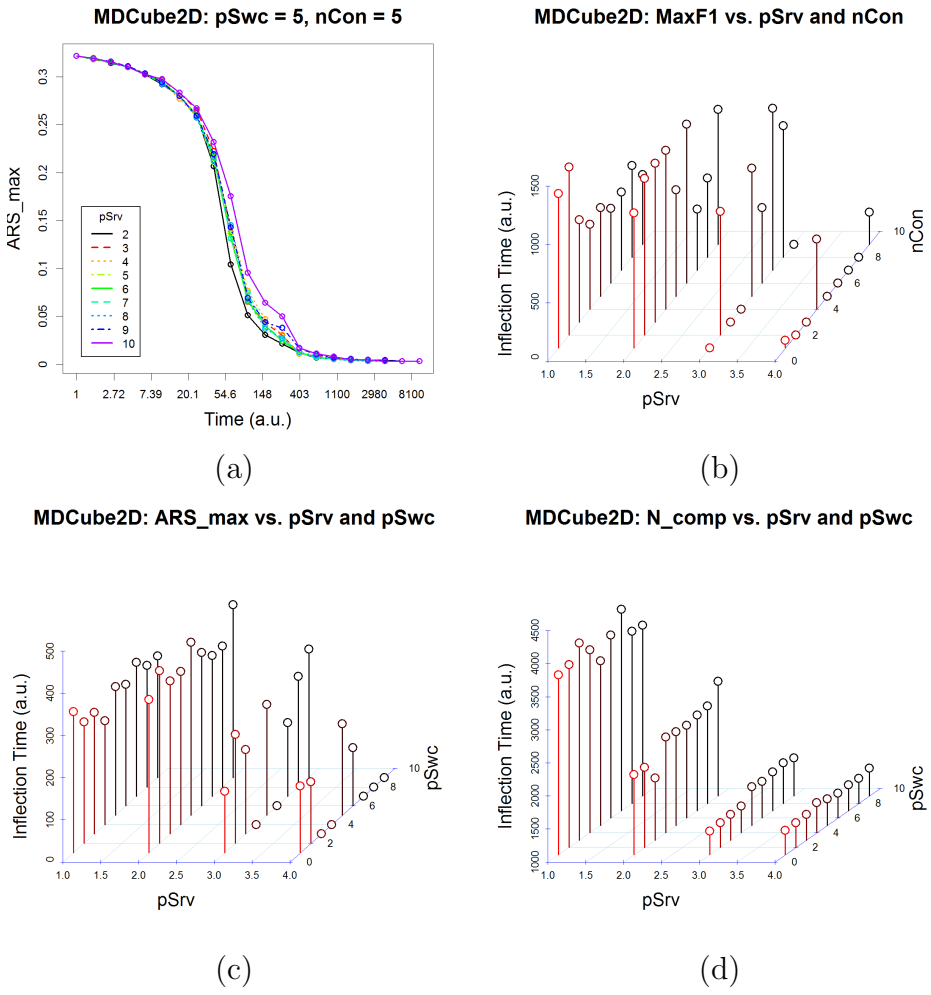

Figure 13: a) The profile of the ARS_{max} metric for the MDCube2D topology. b)-d) The variation of $Flow_{max}$, ARS_{max}, and N_{comp} metrics versus the MDCube2D topology's parameters.

performed. In this section, the performance comparison is performed for the same number of 3,072 servers in each topologies. The original BCube, FatTree, and MDCube2D topologies require a large number of units in order to realize this number of servers.[18] For example, using 48-port switches, a 2304-server BCube (level-1) architecture can be achieved. The next DC size would be 2704 using 52-port switches. Similarly, the FatTree topology could achieve the same number of servers using these switches. In contrast, our proposed BCube* topology can provide an level-1 architecture using 64 BCube (level-0) cubes each one containing 48 servers to easily archive the target 3072-server size. As designed in the definition of the BCube* topology, The level-1 48-port switches contribute their ports to create "imaginary" switches of 64 ports that handle the traffic from the 64 level-0 cubes at the level one. In this section,

[18] Alternatively, the number of switches' port can made variable across the levels. However, in this work, we suggest single-type switch designs in order to reduce the sensitivity to hardware and also avoid specialized and more expensive hardware.

Topology	pSrv	pSwc	# servers	# switches
BCube*	2	48	3072	128
FatTree*	2	48	3072	134
MDCube2D*	2	48	3072	128

Table 1: The parameters and specifications of the three MDC topologies used for the 3072-server use case.

for the purpose of comparability across the three topologies, we do not consider higher levels in the BCube* topology. However, in practice, BCube* architectures with higher levels of connectivity can be considered with a single type of switches. Also, it is worth noting that the resolution in increasing or decreasing the DC size is the size of one level-0 cube, i.e., 48 servers, for the BCube* topology. It is 400 (=2704-2304) in the case of BCube topology when replacing 48-port switches with 52-port ones. A comparison table of the required number of servers and switches of the three extended topologies is provided in Table 1. In particular, the MDCube2D* architecture composed of four sub BCube* architectures each one containing 768 servers. The FatTree* architecture composed of 64 pods, and the BCube* architecture has 64 level-0 cubes.

First, the performance of the BCube* design is analyzed in Figure 14. Five metrics, ARS_{max}, RS_{max}, ARS_{Tot}, $MaxF_1$, and $MaxF_2$, are considered. As can be seen, the RS_{max} measure is not a highly indicative measure because it ignores the actual power of the surviving components and pools. Also, ARS_{Tot} closely follows the ARS_{max} that is a good performance for the BCube topologies in profitability of the solution by providing service to multiple customer or tenants along the time. The resilience to breakdown to multiple pools can be seen from the $MaxF_1$ and $MaxF_2$ performance shown in Figure 14(b). Considering the 2-port servers included in the design, the upper limit of the $MaxF_1$ is 2, and it degrades along the time toward 1. The $MaxF_1$ does not go below one because of its definition that calculates the flow within a pool and island, which in turn implicitly enforces minimal connectivity in order to have a connect component.

The same analysis, but for the MDCube2D* topology, is presented in Figure 16. As expected the breakdown moves to earlier time periods. Also, the performance of the ARS_{Tot} is degraded considerably. Finally, the performance of the FatTree* design is shown in Figure 15. Interestingly, the breakdown time period of the FatTree* architecture does not match the period of degradation associated with the biggest pool (as can be seen from Figure 15(b), the breakdown is around time equal to 80, while the biggest pool degrades to a flow of one around time equal to 700). This can be explained by multiple partitioning events that happen at the breakdown time and result in several small pools. In contrast, the BCube* and the MDCube2D* topologies, observe a partitioning into a big and small pool at the breakdown.

In order to have a comparative analysis, the performance curves of the three designs are superimposed and shown in Figure 17. In support of previously mentioned points, the BCube*

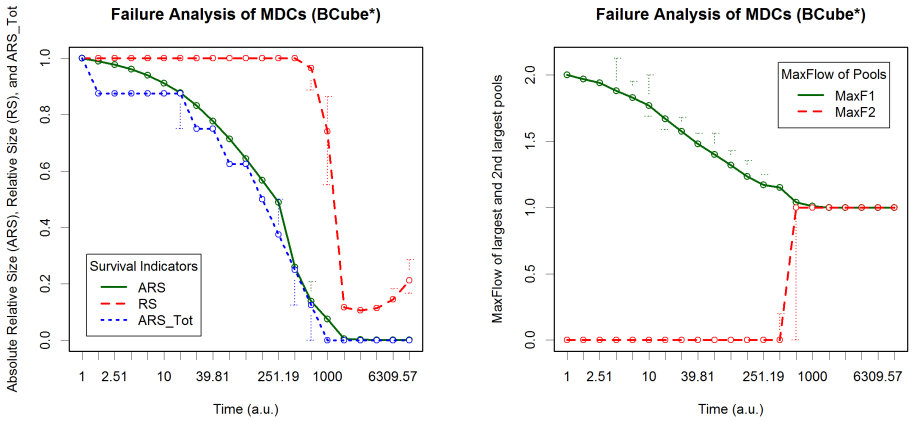

(a) (b)

Figure 14: The time profile of failure performance of the BCube* topology versus various size and flow metrics.

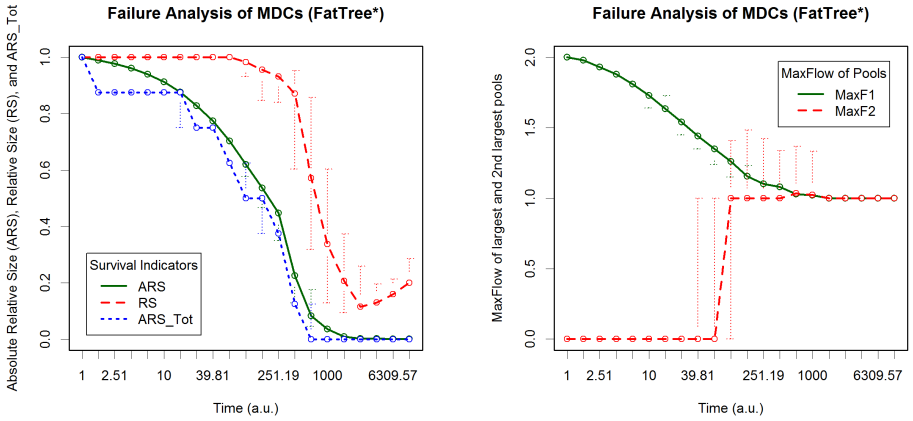

(a) (b)

Figure 15: The time profile of failure performance of the FatTree* topology versus various size and flow metrics.

topology performs the best in terms of all three measure of the ARS_{max}, ARS_{Tot}, and $MaxF_1$. On the second rand, the MDCube2D* topology stands with an exception with respect to the $MaxF_1$. The bottleneck of the MDCube2D* topology is in its few number of inter-container links that can easily downgrade its flow performance. Considering the goal of this paper, which is designing architectures for MDCs to be deployed in zero- or low-maintenance places, this could impose a challenge. However, for those cases that a minimal access to the site is provisioned, the degradation can be easily restored. All these insights are calculated thanks to the proposed lifetime analysis approach that allows calculation of the performance metrics along time, and not at a selected probability rate.

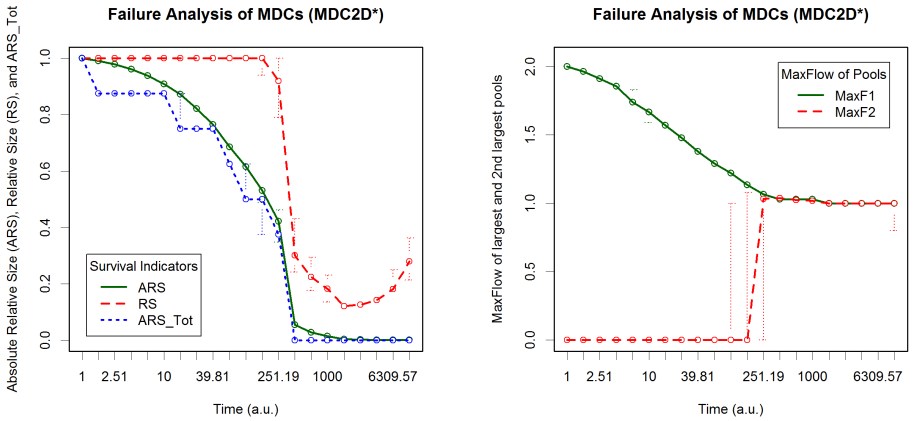

(a) (b)

Figure 16: The time profile of failure performance of the MDCube2D* topology versus various size and flow metrics.

8. Related Work

It is always important to use a sophisticated methodology to provide a comparison study on the reliability and scalability in DC network topologies. There are several work addressing the reliability concern in DC networks [11, 23, 24, 40, 51]. In these studies, a few DC network architectures were investigated in order to provide a better reliability and scalability with MDCs. In this regard, a comparison study on the reliability of DC topologies such as FatTree, BCube and DCell, were investigated in [23, 24] subject to a specific traffic pattern and routing algorithm, respectively. In particular, it has been shown that BCube topology that is specifically designed for MDCs using shipping-containers, shows a better performance than FatTree and DCell topologies for MDCs. Moreover, there are number of studies those taken into account comparisons based on the energy conservation and cost of different DC network architectures [40, 51]. For instance, the authors in [40] have taken one step to understand the trade-offs in the design of DC network architectures. In particular, they proposed a methodology to estimate the cost of various DC architectures based on a predetermined performance requirement. Then, making use of the proposed methodology, they evaluated different DC network topologies and compared them in terms of the cost and power consumption aspects. On the other hand, another step have taken in [51] in which the authors proposed a methodology to analyze and compare the energy consumption in various DC network architectures. Specifically, the impact of considering an energy-aware routing protocol on the energy consumption of DC networks has been investigated in [51] for various DC network topologies. It has been shown that the BCube topology with the proper energy-aware routing protocol has a better energy consumption than FatTree and DCell topologies.

Only recently, independent of any network applications, the authors in [11] investigated the reliability of different DC network topologies in terms of the network components' failure ratios.

In these analysis, the DC network topologies were modeled as a graph with servers and switches as nodes and network links that connecting the nodes. The results has shown a degradation at the DC network performance because of the removal of connected components with a small number of servers as the failure ratio increases. However, please note that the reliability comparisons analysis provided in studies like [11] were utilized the removal of different network components to evaluate the effect of each component's failure on the reliability of different DC network topologies. It is worth noting that though such a comparison is common and used in many contributions, it is also challenging since it does not consider the failure behavior of each component over time. In fact, up to our knowledge, there exists no study on comparing different DC network topologies in terms of failures on DC network components while taking into account the operational life-time (lifespan) of the network.

9. Conclusion and future prospects

In this work, a novel statistical approach, based on the Monte-Carlo methodology, is proposed to estimate the performance metrics of modular data centers along their operating lifespan. To our knowledge, this is the first time such an analysis is provided. The proposed approach breaks down the trajectory dependency of the system in order to develop an unbiased statistical picture of a design performance along its utilization. The approach uses the components probability distributions of failure along their utilization time in order to calculate the failure penetration at the components level, referred to as a snapshot. At the same time, to have a better fit on the real system failure data, a generalization to the Tanh-Log cumulative probability distribution has been proposed. Using the proposed distribution and the analysis approach, the performance of three topologies in the context of modular data centers has been studied. Also, in order to make these topologies more flexible with respect to the hardware and increase their resilience to failures, some extended topologies have been introduced. It has been concluded that the extended BCube topology delivers the highest resiliency performance in terms of biggest pool size, total pool size, and flow volume. With respect to scalability, the extended MDCube2D provides the best performance. Study of the impact of other layers and components, such as virtual machine managers, virtual machines, and in general software on the reliability, availability, and management of data centers will be considered in the future.

Acknowledgments

The authors thank the NSERC of Canada for their financial support under grant CRDPJ 424371-11.

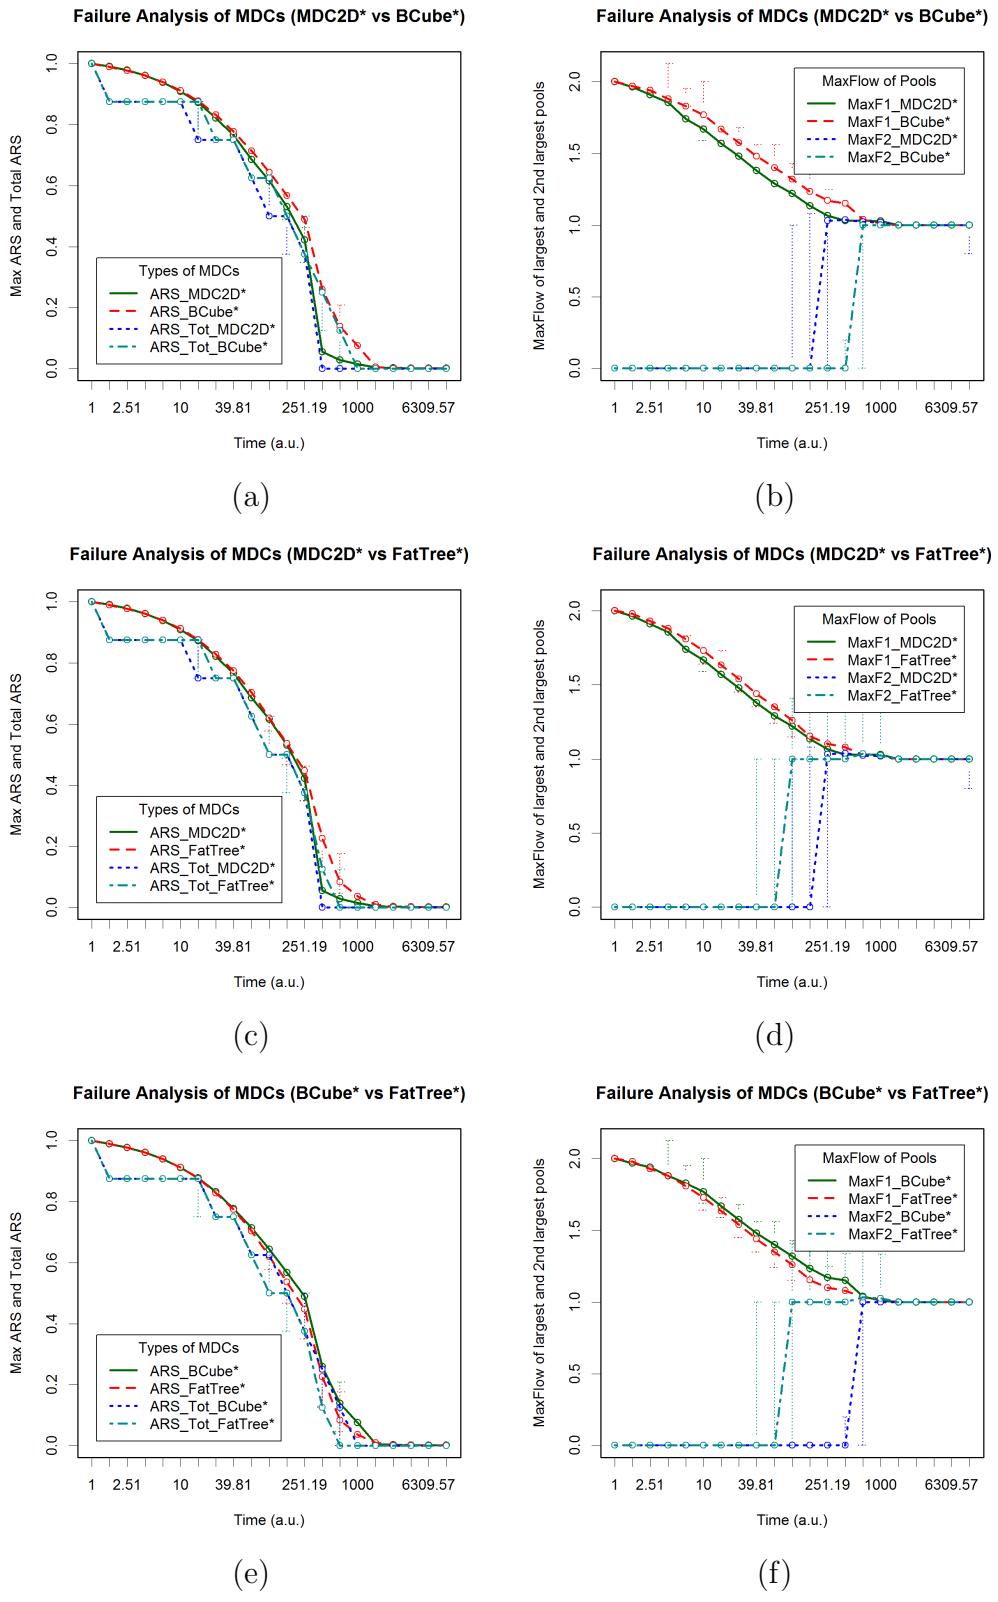

Figure 17: The pairwise comparison of the three proposed topologies against size and flow metrics. a)-b) MDCube2D* vs BCube*. c)-d) MDCube2D* vs FatTree*. e)-f) BCube* vs FatTree*.

References

[1] W.A. Abdelmaksoud, T.Q. Dang, H.E. Khalifa, R.R. Schmidt, and M. Iyengar. Perforated tile models for improving data center CFD simulation. In *ITherm'12*, pages 60–67, San Diego, CA, USA, May 30 - June 1 2012.

[2] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, July 2000.

[3] Luiz André Barroso and Urs Hölzle. The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis Lectures on Computer Architecture*, 4(1):1–108, January 2009.

[4] T Berger, KA Steven, and H King. *Canada's North: What's the plan?* The 2010 CIBC scholar-in-residence lecture. The Conference Board of Canada, Ottawa, Ontario, Canada, 2010.

[5] Levi A Campbell, Richard C Chu, Michael J Ellsworth Jr, Madhusudan K Iyengar, and Robert E Simons. Liquid-cooled electronics rack with immersion-cooled electronic subsystems. U.S. Patent 8,184,436, Filed on Jun 29, 2010, Published on May 22, 2012.

[6] Ramkumar Chinchani, Duc Ha, Anusha Iyer, Hung Q. Ngo, and Shambhu Upadhyaya. Insider threat assessment: Model, analysis and tool. In Scott C.-H. Huang, David MacCallum, and Ding-Zhu Du, editors, *Network Security*, pages 143–174–. Springer US, 2010.

[7] Minki Cho, N. Sathe, A. Raychowdhury, and S. Mukhopadhyay. Optimization of burn-in test for many-core processors through adaptive spatiotemporal power migration. In *ITC'10*, pages 1–9, 2010.

[8] Demetris T Christopoulos. Developing methods for identifying the inflection point of a convex/concave curve. *arXiv preprint arXiv:1206.5478*, June 24 2012.

[9] R.C. Chu, R.E. Simons, M.J. Ellsworth, R.R. Schmidt, and V. Cozzolino. Review of cooling technologies for computer products. *IEEE Transactions on Device and Materials Reliability*, 4(4):568–585, 2004.

[10] C. Clos. A study of nonblocking switching networks. In *Bell System Technical Journal*, pages 406–424, 1953.

[11] R.S. Couto, M.E.M. Campista, and L.H.M.K. Costa. A reliability analysis of datacenter topologies. In *GLOBECOM'12*, pages 1890–1895, Anaheim, CA, USA, Dec 3-7 2012.

[12] Jeffrey Dean and Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. In *OSDI'04*, San Francisco, CA, USA, Dec 6-8 2004.

[13] J.Mason Earles and Anthony Halog. Consequential life cycle assessment: a review. *The International Journal of Life Cycle Assessment*, 16(5):445–453, 2011.

[14] Nosayba El-Sayed, Ioan A. Stefanovici, George Amvrosiadis, Andy A. Hwang, and Bianca Schroeder. Temperature management in data centers: why some (might) like it hot. *SIGMETRICS Perform. Eval. Rev.*, 40(1):163–174, 2012.

[15] Nosayba El-Sayed, Ioan A. Stefanovici, George Amvrosiadis, Andy A. Hwang, and Bianca Schroeder. Temperature management in data centers: why some (might) like it hot. In *IGMETRICS'12*, pages 163–174, London, England, UK, June 11-15 2012.

[16] Fereydoun Farrahi Moghaddam, Mohamed Cheriet, and Kim Khoa Nguyen. Low carbon virtual private clouds. In *CLOUD'11*, pages 259–266, Washington, DC, USA, 4-9 July 2011.

[17] Fereydoun Farrahi Moghaddam, Reza Farrahi Moghaddam, and Mohamed Cheriet. Carbon metering and effective tax cost modeling for virtual machines. In *CLOUD'12*, pages 758–763, Honolulu, Hawaii, USA, June 2012.

[18] Reza Farrahi Moghaddam, Fereydoun Farrahi Moghaddam, Vahid Asghari, and Mohamed Cheriet. Cognitive behavior analysis framework for fault prediction in cloud computing. In *NoF'12*, pages 1–8, Tunis, Tunisia, Nov 21-23 2012.

[19] Adam Fiser. Mapping the long-term options for Canada's North: Telecommunications and broadband connectivity. Technical report, The Conference Board of Canada, August 2013.

[20] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(35):75–174, February 2010.

[21] Stefan Fournier. Getting it right: Assessing and building resilience in Canada's North. National security and public safety, The Conference Board of Canada, May 2012.

[22] Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, Parveen Patel, and Sudipta Sengupta. VL2: a scalable and flexible data center network. *SIGCOMM Comput. Commun. Rev.*, 39(4):51–62, 2009.

[23] C. Guo, G. Lu, D. Li, H. Wu, X. Zhang, Y. Shi, C. Tian, Y. Zhang, and S. Lu. BCube: a high performance, server-centric network architecture for modular data centers. In *SIGCOMM'09*, pages 63–74, Barcelona, Spain, August 17-21 2009.

[24] C. Guo, H. Wu, K. Tan, L. Shi, Y. Zhang, and S. Lu. DCell: a scalable and fault-tolerant network structure for data centers. In *SIGCOMM'08*, pages 75–86, Seattle, WA, USA, August 17-22 2008.

[25] J. Hamilton. Architecture for modular data centers. *eprint arXiv:cs/0612110*, December 2006.

[26] Karen Hofman and Xianguo Li. Canada's energy perspectives and policies for sustainable development. *Applied Energy*, 86(4):407–415, April 2009.

[27] Andrey N. Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *Giornalle dell'Instituto Italiano degli Attuari*, 4(1):83–91, 1933.

[28] Derrick Kondo, Bahman Javadi, Alexandru Iosup, and Dick Epema. The failure trace archive: Enabling comparative analysis of failures in diverse distributed systems. In *CCGrid'10*, pages 398–407, Melbourne, VIC, Australia, May 17-20 2010.

[29] Antonios Litke, Dimitrios Skoutas, Konstantinos Tserpes, and Theodora Varvarigou. Efficient task replication and management for adaptive fault tolerance in mobile grid environments. *Future Generation Computer Systems*, 23(2):163–178, February 2007.

[30] M. Al-Fares, A. Loukissas and A. Vahdat. A scalable, commodity data center network architecture. *SIGCOMM Comput. Commun. Rev.*, 38(4):63–74, 2008.

[31] Jackson Braz Marcinichen, John Richard Thome, and Bruno Michel. Cooling of microprocessors with micro-evaporation: A novel two-phase cooling cycle. *International Journal of Refrigeration*, 33(7):1264–1276, November 2010.

[32] Frank J. Massey. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American Statistical Association*, 46(253):68–78, March 1951.

[33] Nick McKeown. Software-defined networking. In *INFOCOM'09 keynote talk*, April 2009. Available: http://tiny-tera.stanford. edu/ nickm/talks/infocom_brazil_2009_v1-1.pdf.

[34] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. OpenFlow: enabling innovation in campus networks. *SIGCOMM Comput. Commun. Rev.*, 38(2):69–74, 2008.

[35] Matthew Monaco, Oliver Michel, and Eric Keller. Applying operating system principles to SDN controller design. In *HotNets'13*, pages 1–7, College Park, Maryland, Nov 21–22 2013. ACM.

[36] Javier Navaridas, Jose Miguel-Alonso, Francisco Javier Ridruejo, and Wolfgang Denzel. Reducing complexity in tree-like computer interconnection networks. *Parallel Comput.*, 36(2-3):71–85, 2010.

[37] P. Fuentes, E. Vallejo, C. Martinez, M. Garcia and R. Beivide. Comparison study of scalable and cost-effective interconnection networks for HPC. In *ICPPW'12*, pages 594–595, 2012.

[38] Martin Pehnt, Michael Oeser, and Derk J. Swider. Consequential environmental system analysis of expected offshore wind electricity production in Germany. *Energy*, 33(5):747–759, May 2008.

[39] Fabrizio Petrini and Marco Vanneschi. k-ary n-trees: High performance networks for massively parallel architectures. In *IPPS'97*, pages 87–93, Geneva, Switzerland, April 1-5 1997.

[40] L. Popa, S. Ratnasamy, G. Iannaccone, A. Krishnamurthy, and I. Stoica. A cost comparison of datacenter network architectures. In *CoNEXT10*, pages 16:1–16:12, Philadelphia, PA, USA, Nov 30 - Dec 3 2010.

[41] M. Ramfitt and H. Coles. Modular/container data centers procurement guide: Optimizing for energy efficiency and quick deployment. Technical report, Lawrence Berkeley National Laboratory on behalf of the General Services Administration (GSA), February 2 2011. Available at http://goo.gl/ho9cC.

[42] Kim Nguyen Yves Lemieux Réjean Samson Mohamed Cheriet Reza Farrahi Moghaddam, Thomas Dandres. A cross-analysis of ICT life-cycle assessment standards & application in green sustainable Telco cloud. In *SETAC'14*, page Submitted, Basel, Switzerland, May 11-15 2014.

[43] Roberto Riggio, Tinku Rasheed, and Fabrizio Granelli. EmPOWER: A testbed for network function virtualization research and experimentation. In *SDN4FNS'13*, pages 1–5, Trento, Italy, Nov 11-13 2013.

[44] B. Schroeder and G.A. Gibson. A large-scale study of failures in high-performance computing systems. *IEEE Transactions on Dependable and Secure Computing*, 7(4):337–350, 2010.

[45] Michael A Stephens. Goodness of fit tests with special reference to tests for exponentiality. Technical Report 262, Department of Statistics, Stanford University, Stanford, CA, USA, Sept 18 1978.

[46] Chih-Chun Tsai, Sheng-Tsaing Tseng, and N. Balakrishnan. Optimal burn-in policy for highly reliable products using Gamma degradation process. *IEEE Transactions on Reliability*, 60(1):234–245, 2011.

[47] Chih-Chun Tsai, Sheng-Tsaing Tseng, and N. Balakrishnan. Optimal design for degradation tests based on Gamma processes with random effects. *IEEE Transactions on Reliability*, 61(2):604–613, 2012.

[48] Balajee Vamanan, Jahangir Hasan, and T.N. Vijaykumar. Deadline-aware datacenter TCP (D2TCP). In *SIGCOMM'12*, pages 115–126, Helsinki, Finland, August 13-17 2012. ACM.

[49] Haitao Wu, Guohan Lu, Dan Li, Chuanxiong Guo, and Yongguang Zhang. MDCube: a high performance network structure for modular data center interconnection. In *CoNEXT'09*, CoNEXT'09, pages 25–36. ACM, 2009.

[50] Shaomin Wu and Min Xie. Classifying weak, and strong components using ROC analysis with application to burn-in. *IEEE Transactions on Reliability*, 56(3):552–561, 2007.

[51] Y. Shang, D. Li and M. Xu. A comparison study of energy proportionality of data center network architectures. In *ICDCSW'12*, pages 1–7, Macau, China, June 18-21 2012.

[52] David Zats, Tathagata Das, Prashanth Mohan, Dhruba Borthakur, and Randy Katz. DeTail: reducing the flow completion time tail in datacenter networks. *SIGCOMM Comput. Commun. Rev.*, 42(4):139–150, 2012.

[53] Xuanhang (Simon) Zhang, J.W. VanGilder, M. Iyengar, and R.R. Schmidt. Effect of rack modeling detail on the numerical results of a data center test cell. In *ITherm'08*, pages 1183–1190, Orlando, FL, USA, May 28-31 2008.

Reza FARRAHI MOGHADDAM received his B.Sc. degree in Electrical Engineering and his Ph.D. degree in Physics from the Shahid Bahonar University of Kerman, Iran, in 1995 and 2003, respectively. He has been a Postdoctoral Research Fellow and a Research Associate with the Synchromedia Laboratory for Multimedia Communication in Telepresence, École de technologie supérieure (University of Quebec), Montreal (QC), Canada since 2007 and 2012, respectively. Dr. Farrahi has published more than 50 technical papers, and his research interests include green computing, sustainable ICT, behavior analysis, life cycle analysis, green economy, image processing, visual perception, and optimization.

Vahid ASGHARI obtained his Ph.D. degree in Telecommunications from INRS-EMT, University of Quebec, Montreal, Canada in 2012 and his B.Sc. degree in electrical engineering from Azad University and the M.Sc. degree in telecommunication systems from K.N. Toosi University of Technology in 2002 and 2005, respectively. He is currently working as a postdoctoral research fellow at McGill University, Canada. His research interests include resource scheduling and management, cooperative communications with a focus on heterogeneous systems. He is also interested in behavior modeling, analysis and management of smart applications using ICT technology with a focus on cloud computing networks. He is a recipient of Postdoctoral Research Fellowship Award from the Quebec Government FQRNT, 2012. He is also co-recipient of Best Paper Award from IEEE WCNC 2010.

Fereydoun FARRAHI MOGHADDAM received his B.Sc. degree in Electronics Engineering from Shahid Bahonar University of Kerman, Iran, in 1999. He obtained his M.Sc. degree in Electronics Engineering from K.N. Toosi University of Technology, Tehran, Iran, in 2001. He joined the Synchromedia Laboratory for Multimedia Communication in Telepresence, École de technologie supérieure (University of Quebec), Montreal (QC), Canada as a Ph.D. student in September 2008, and received his Ph.D. degree on green distributed networks of data centers in 2014. He is currently working as a Postdoctoral Research Fellow with Synchromedia Lab since 2014. His research interests include green cloud computing, smart scheduling, and optimization.

Mohamed CHERIET received his B.Eng. from USTHB University (Algiers) in 1984 and his M.Sc. and Ph.D. degrees in Computer Science from the University of Pierre et Marie Curie (Paris VI) in 1985 and 1988 respectively. Since 1992, he has been a professor in the Automation Engineering department at the École de technologie supérieure (University of Quebec), Montreal, and was appointed full professor there in 1998. His interests include document image analysis, OCR, mathematical models for image processing, pattern classification models and learning algorithms, as well as perception in computer vision. Dr. Cheriet has published more than 250 technical papers in the field, and has served as chair or co-chair of the following international conferences: VI'1998, VI'2000, IWFHR'2002, ICFHR'2008, and ISSPA'2012.